

The Removed First Rung: Digital Transformation and the Reconstruction of Talent Architecture in IT Organisations

AUTHOR:- Mrs.SHAILA ASHOK RANE

RESEARCH STUDENT

BHARATI VIDYAPEETH DEEMED UNIVERSITY, PUNE

Email Id: Shailarane780@gmail.Com

UNDER THE SUPERVISION OF

Dr.VINOD H. MANE (MA, M.Phil, G.D.C.& A., Ph.D.)

Abstract

Talent management rests on an assumption that is rarely stated: that the positions being managed hold still long enough to be identified, ranked, and filled. In IT organisations that assumption is now under pressure from an unusual direction. The technology these firms adopt does not merely serve the talent system; it acts on the work from which the talent categories are built. When routine implementation becomes tool-assisted, the entry-level engineering role does not simply contract in headcount. It changes in kind, and the developmental sequence that converted junior hires into senior technical capability loses the step it began with. This paper sets out a research design for studying that problem, and distinguishes it from the more familiar argument about skills obsolescence: obsolescence concerns the content of competence, whereas the concern here is the structural mechanism by which competence is produced. Four propositions are developed, linking digital transformation intensity to sourcing criteria, entry-portal width, the mutual constitution of talent-system response and pipeline architecture, and the moderating role of contractual form. The design is interpretivist across three sequential phases and configurational in a fourth: document and archival analysis of approximately 40 artefacts, 30 semi-structured interviews stratified across three role strata and three firm types, 60 to 70 critical incidents collected using Flanagan's technique, and a fuzzy-set qualitative comparative analysis treating firms as configurations of five conditions. Coding follows the Gioia methodology with a pre-declared saturation criterion and reported intercoder reliability. No findings are presented. What is presented is the apparatus, specified in enough detail to be assessed, replicated, or contested before fieldwork begins.

Keywords: talent management; digital transformation; capability pipeline; entry-level automation; fuzzy-set qualitative comparative analysis; IT workforce; fissured workplace

1. The research problem

A software engineer joining an IT organisation in 2019 usually began with work that someone more senior had already specified: a defined function, a bounded defect, a failing test to make pass. The work was small, and the smallness was the point. It was reviewable, which made corrective feedback possible; it was low-stakes, which made errors affordable; and it exposed the newcomer to the codebase, the deployment path, and the review conventions of the team without requiring anyone to teach those things explicitly. Competence was produced as a by-product of production. That arrangement, which most IT firms operated without ever writing it down, is what tool-assisted code generation has disturbed.

Talent management as a field grew from a scarcity premise. Collings and Mellahi (2009) gave it its most durable formulation: identify the pivotal positions that differentially affect strategic outcomes, fill them with high-potential incumbents, and build the architecture that develops and retains those incumbents. The formulation has been productive

and it has also been criticised, most directly by Lewis and Heckman (2006), who argued that the field had assembled a vocabulary faster than it had assembled a construct. Later work sharpened the object without disturbing its foundation. Gallardo-Gallardo, Dries and Gonzalez-Cruz (2013) asked what talent actually denotes; Sparrow and Makram (2015) asked where its value is created; Vaiman, Scullion and Collings (2012) examined how talent decisions get made in practice; Boudreau and Ramstad (2005) proposed a decision science for differentiated investment. Across all of it one structural assumption persists. Positions exist prior to the talent system, and the system's work is to fill them.

Digital transformation research proceeds from a different assumption. Vial's (2019) synthesis treats transformation as a process in which digital technologies trigger changes to value creation paths, organisational structures, and the barriers that constrain both, rather than as a procurement decision with an implementation timeline attached. Warner and Wäger (2019) read the same process through dynamic capabilities, describing the sensing, seizing, and transforming routines through which incumbents rebuild themselves, and drawing on Teece's (2007) account of how those routines operate. Neither body of work says much about what transformation does to the positions a talent system is meant to fill. In most sectors that silence is reasonable, because the technology being adopted and the work that produces senior capability are separated by several organisational layers. A bank deploying a customer analytics platform is not thereby altering how its relationship managers learn their trade.

IT organisations are the exception, and that is what makes them the right setting rather than merely a convenient one. Here the technology being adopted acts directly on the work of the people adopting it. When routine implementation becomes tool-assisted, the effect on the entry-level role is not confined to headcount. The tasks that constituted the role are absorbed, and what remains under that job title is a different activity: reviewing generated output, specifying problems precisely enough for a tool to act on them, deciding when a plausible answer is wrong. The talent architecture built on the old role no longer describes anything. Neither literature has examined what happens to a talent system when the technology it was built to serve dissolves the developmental sequence the system assumes.

That is the gap this study occupies, and it is worth separating immediately from an argument it resembles. Skills obsolescence concerns the content of competence: which languages, frameworks, and platforms a worker needs, and how quickly that inventory turns over. Firms have managed content turnover for four decades through retraining, certification, and lateral hiring, and the management literature has followed them. The claim advanced here concerns the mechanism by which competence is produced rather than its content. A curriculum can be rewritten. A developmental sequence cannot be rewritten in the same way, because it was never a curriculum; it was a by-product of how work happened to be divided, and when the division changes the by-product stops appearing. Obsolescence is a problem of updating. This is a problem of reproduction.

Four questions follow from that framing.

RQ1. How do IT organisations reconfigure sourcing, development, and retention practices when routine technical tasks become automatable?

RQ2. Through what mechanisms does the automation of entry-level work alter the progression architecture that produces senior technical capability?

RQ3. Which combinations of organisational conditions are sufficient for maintaining capability depth under high digital transformation intensity?

RQ4. How does the contractual form of the technical workforce condition the relations in RQ1 to RQ3?

RQ3 is phrased in the language of set-theoretic sufficiency deliberately. It is answered configurationally in Phase 4, on the expectation that no single condition produces capability depth on its own, and the phrasing is meant to make that expectation visible rather than to smuggle it in at the analysis stage.

This study contributes in four ways. First, it inverts the direction of the technology–talent relation that talent management research has assumed, treating the technology as something that reorganises the categories the talent system operates on rather than as an instrument the system deploys. Second, it specifies a mechanism, the removed first rung, that is observable in artefacts firms already produce and archive: job architecture documents, competency frameworks, campus hiring materials, and internal mobility policies, together with their revision histories. Third, it brings Flanagan's (1954) critical incident technique into a literature that has relied heavily on retrospective accounts of policy, on the grounds that

specific remembered events are harder to narrate into compliance with official practice than general questions are. Fourth, it treats the contractual form of the technical workforce as a boundary condition rather than as a control variable, following Weil (2014), which allows the design to speak to a workforce composition that is characteristic of the sector and largely absent from the talent management literature.

One matter of scope should be stated plainly at the outset. This is a design paper. No findings are reported, because fieldwork has not been conducted. What is reported is the analytical apparatus in enough detail that it can be assessed, replicated, or contested before data collection commits the study to a particular set of choices.

2. Theoretical framework

2.1 What talent management assumes

The strategic talent management position holds that differentiated investment in a small set of positions yields performance advantage disproportionate to its cost. Collings and Mellahi (2009) locate the differentiation at the level of the position rather than the individual, which is what distinguishes their formulation from earlier succession-planning approaches, and Boudreau and Ramstad (2005) supply the decision logic for where marginal investment should go. Cappelli (2008) added a supply-chain reading, arguing that firms should manage talent under uncertainty rather than plan against forecasts they cannot support. Each of these arguments presupposes that the positions can be identified. Identification requires stability over the planning horizon, and the horizon in question is long: the interval between hiring a graduate engineer and that person becoming a senior technical contributor runs to five or seven years in most IT organisations. A talent architecture is therefore a bet on the persistence of a role structure across an interval during which the technology of the work may change entirely.

2.2 Human capital, capability depth, and reconfiguration

Becker's (1964) distinction between general and firm-specific human capital explains why firms train at all. A firm captures the return on training only to the extent that it retains the trained worker, which makes investment in general skills rational for the worker and marginal for the employer, and investment in firm-specific skills rational for the employer to the extent that tenure is expected. Apprenticeship arrangements in engineering teams survive on that logic. They also depend on something Becker's framework does not address, because it takes the training task as given: they require the existence of work that is simultaneously productive and instructive. Remove that work and the calculation does not become unfavourable, it becomes inapplicable.

The resource-based view supplies the reason to care. Barney (1991) identifies sustained advantage with resources that are valuable, rare, imperfectly imitable, and non-substitutable, and deep senior technical capability qualifies on all four counts, particularly in firms whose competitive position rests on systems that cannot be reconstructed from documentation. What the resource-based view does poorly is explain how such a resource forms over time, a limitation the dynamic capabilities literature was developed to address. Teece (2007) supplies the vocabulary of sensing, seizing, and reconfiguring, and Warner and Wäger (2019) apply it to digital transformation directly. The vocabulary operates at the level of the firm. It says little about the individual progression sequences through which the reconfigured firm would obtain the people it needs.

2.3 Individual-level theories and the contractual boundary

Two individual-level theories connect the structural argument to observable outcomes. Rousseau's (1995) psychological contract predicts that unwritten reciprocal obligations, once broken, produce responses out of proportion to the material loss involved, which matters here because progression promises are among the most common unwritten obligations in technical employment. An engineer who joined on an implicit understanding about what the second and third years would contain, and who finds those years occupied by reviewing generated output, has experienced a breach whether or not any policy was violated. Job embeddedness (Mitchell, Holtom, Lee, Sablinski, and Erez, 2001) predicts that links, fit, and sacrifice retain employees beyond what satisfaction alone would explain, and it predicts differential turnover across contractual forms, since contractors accumulate fewer links and forfeit less by leaving. The theory was developed

before distributed work became standard in the sector, and its boundary conditions under those arrangements are not well established.

Weil's (2014) account of the fissured workplace supplies the final input. Work has been displaced from the firms that direct and benefit from it to vendors, staffing intermediaries, and capability centres operating under service agreements, and the IT sector is among the most heavily fissured. The consequence for the present argument is direct. A progression architecture is an instrument of the employing entity, and it governs only those inside the employment perimeter. Table 1 states what each theory contributes and what it does not.

Table 1. Theoretical inputs and what each contributes

Theory	Core proposition	Contribution here	What it does not supply
Strategic talent management (Collings & Mellahi, 2009)	Differentiated investment in pivotal positions produces performance advantage	Defines the object being reconfigured	Assumes positions are stable enough to be identified
Human capital theory (Becker, 1964)	Firms underinvest in general skills and overinvest in firm-specific ones	Explains why apprenticeship is firm-rational only when tenure is expected	Silent on what happens when the training task disappears
Resource-based view (Barney, 1991)	Rare, inimitable, non-substitutable resources yield sustained advantage	Frames capability depth as the asset at risk	Static; poor at explaining resource formation over time
Dynamic capabilities (Teece, 2007)	Sensing, seizing, and reconfiguring drive adaptation	Supplies the reconfiguration vocabulary	Operates at firm level; little on individual progression
Psychological contract (Rousseau, 1995)	Employment involves unwritten reciprocal obligations	Explains retention effects when progression promises break	Individual level; needs linking to structure
Job embeddedness (Mitchell et al., 2001)	Links, fit, and sacrifice retain employees beyond satisfaction	Predicts differential turnover across contractual forms	Developed pre-remote-work; boundary conditions unclear
Fissured workplace (Weil, 2014)	Work is displaced from directing firms to vendors and intermediaries	Supplies the contractual boundary condition in P4	Descriptive of structure; not of capability formation

2.4 Proposition development

Selection instruments encode a theory of what is scarce. When implementation was the constraint, firms tested implementation: a whiteboard algorithm, a timed coding screen, a take-home exercise judged on correctness and style. These instruments discriminated well because writing correct code under constraint was difficult and, more importantly, because the difficulty was predictive of the difficulty of the job. Tool-assisted generation compresses that difficulty

substantially. It compresses far less of the work of deciding what should be built, recognising output that is wrong in ways tests do not catch, and specifying a problem precisely enough that a generated answer is usable. If the scarce input has moved, the instruments that identify it have reason to move as well, and role definitions should follow, since job architecture documents are where firms record what they believe they are buying.

P1. Higher digital transformation intensity produces a shift in sourcing criteria from demonstrated technical execution toward judgement, review, and specification capability, observable in role definitions and selection instruments; that is, the more pervasive tool-assisted execution becomes within a firm, the more its hiring documentation and selection design emphasise evaluation and specification over demonstrated implementation.

Automation does not distribute evenly across a hierarchy. It concentrates where work is most codified, most repetitive, and most amenable to review against a specification, which describes entry-level engineering reasonably well and describes architecture decisions, incident command, and scope negotiation with a business owner rather poorly. The consequence is not proportional contraction across levels. Senior headcount is retained because the work resists substitution; junior headcount absorbs the adjustment, both because the work is more substitutable and because a vacancy at that level is cheaper to leave open through a budget cycle. What results is a structure that looks intact from above and is not. The ladder survives. The step onto it does not.

P2. Automation of routine implementation narrows the entry portal while leaving senior headcount intact, producing a progression architecture with an intact ladder and a removed first rung; that is, graduate and entry-level intake contracts as a share of technical hiring while the level structure above it is retained unchanged.

Competency frameworks are cheap to revise. A human resources function can rewrite one within a quarter, circulate it, and record the revision date, and the exercise consumes little beyond its own drafting time. Rebuilding an entry portal is expensive and slow, because it requires campus relationships, structured mentoring capacity, and senior engineering time redirected from delivery into review of work that would be faster to do directly. Firms facing both tasks under cost pressure will tend to do the cheap one first. The configuration that results, with competence redefined around judgement and the entry portal left narrow, is stable for as long as the existing senior cohort holds, and that cohort was produced by the architecture that has just been dismantled. The deficit is therefore not detectable at the moment it is created. It becomes visible when the trained cohort turns over and the people behind them acquired judgement, if at all, by some route the firm did not design.

P3. Talent-system responses and pipeline architecture are mutually constituting rather than sequential: firms that redefine competence without rebuilding the entry portal generate a capability deficit that surfaces with a lag of several years; that is, the two adjustments determine each other's consequences, and the cost of making only the cheaper one is deferred rather than avoided.

Becker's reasoning applies to the contractual boundary without modification. A firm invests in developing a worker when it expects to capture the return, and expected tenure is what makes that expectation defensible. Where a substantial share of technical labour sits with vendors, staffing intermediaries, or a capability centre operating under a service agreement, three things follow at once. The client firm has no mechanism to develop people it does not employ; the vendor has weak incentive to develop people whose deployment it cannot guarantee beyond the contract term; and the progression architecture, wherever it formally sits, governs a diminishing fraction of the people doing the technical work. Job embeddedness predicts that the same workers accumulate fewer retaining links, which compounds the effect rather than offsetting it.

P4. The effect of P2 on capability depth is conditioned by contractual form: where a substantial share of technical labour sits outside the employment perimeter, progression architecture governs a shrinking fraction of the workforce and the capability deficit is larger; that is, fissuring does not cause the deficit but widens it.




Figure 1. Propositional model. Digital transformation intensity acts on talent-system responses and pipeline architecture, which jointly determine talent outcomes, conditioned by the contractual form of the technical workforce. Schematic; not derived from estimated parameters.

3. Research methodology

3.1 Paradigm and design logic

The study is interpretivist across its first three phases and configurational in the fourth, and the combination requires justification rather than assertion. The object of study is how organisational actors reconstruct categories under uncertainty: what a human resources director now means by junior, what an engineering manager now accepts as evidence of readiness, what an early-career engineer believes the next two years will contain. These are meaning-making processes, and they are not recoverable from measured variables, because any such variable would have to be built from the categories that are themselves in flux. Interviewing, document analysis, and structured incident recall are the appropriate instruments.

RQ3 asks a different kind of question, about which combinations of conditions are sufficient for an outcome across cases, and it needs a technique that compares whole cases rather than isolating the average effect of a condition. Ragin's (2008) fuzzy-set approach does this, and it is more compatible with interpretive work than regression is, because it treats each firm as a configuration rather than as a set of independent observations on separable variables, and because calibration requires the researcher to state substantive thresholds rather than to standardise against a sample distribution. The four phases are sequential in the sense that each supplies inputs to the next, with a feedback path from Phase 2 to sampling. Figure 2 sets out the sequence.




Figure 2. Four-phase research design. Phases 1 to 3 are sequential and interpretive; Phase 4 is configurational and treats the coded firms as cases. The feedback path indicates theoretical sampling: interviewing continues until the saturation criterion in Section 3.5 is met.

3.2 Phase 1: document and archival analysis

Phase 1 produces construct vocabulary rather than findings, and reporting it as anything else would overstate what document analysis can do. Its purpose is to establish how these organisations describe roles in their own language, so that interview questions in Phase 2 use the firm's categories rather than the researcher's imported ones. Asking an engineering manager about the erosion of the entry-level developmental sequence produces a different conversation from asking what the firm's graduate track now involves, and only the second is likely to be answered in terms the respondent uses daily.

The corpus targets approximately 40 documents across 12 to 15 firms: job architecture documents, competency frameworks, campus hiring materials, internal mobility policies, and the workforce sections of annual reports. Analysis follows Bowen (2009), with a structured extraction sheet recording role title, stated competencies, entry requirements, progression criteria, and revision date where the document carries one. The revision dates carry more analytical weight than they might appear to. Because these artefacts are versioned, revision history across successive versions supplies a longitudinal signal without longitudinal fieldwork, and a competency framework rewritten twice within eighteen months is itself evidence about the stability of the categories it defines. Documents are obtained through firm cooperation where possible and from public sources otherwise, and the source is recorded for each, since publicly posted hiring material and internal mobility policy are not equivalent evidence about what a firm does.

3.3 Phase 2: semi-structured interviews

Phase 2 is the core evidentiary phase. Sampling is purposive with maximum variation on two dimensions, firm type and role stratum, giving the matrix in Table 2. The reserve of five interviews is held against the saturation criterion in Section 3.5 rather than allocated in advance.

Table 2. Sampling matrix for Phase 2

Firm type	Talent / HR lead	Engineering manager	Early-career engineer	Subtotal
Product or platform firm	3	3	4	10
Global capability centre	3	3	4	10

Firm type	Talent / HR lead	Engineering manager	Early-career engineer	Subtotal
IT services or vendor firm	3	3	4	10
Total	9	9	12	30 (+5 reserve)

The three role strata are not a robustness measure. They are constitutive of the design, because each stratum has access to a different version of the same object: human resources leads describe intended policy, engineering managers describe enacted practice, and early-career engineers describe experienced reality. Where these accounts diverge, the divergence is a finding about the gap between policy and practice, not measurement error to be reconciled or averaged. A study that interviewed only human resources leads would recover a functioning system almost by construction, since that is the system those respondents are accountable for having built.

The protocol moves from description toward inference across six stages. It opens non-directively, asking the respondent to walk through how a new engineer joined the team three years ago and how that differs now, which establishes the comparison without naming any technology. A task-level probe follows: for a specific recent piece of work, who did which parts, and which of those parts existed as work at all two years ago. Human resources respondents receive a policy probe asking what changed in the competency framework and what triggered the revision. A progression probe asks what someone must now demonstrate to move up, and how they would demonstrate it, which is the question that most often exposes the absence of an answer. A contractual probe asks which of the people doing this work are employees and whether that affects their development. A counterfactual probe closes the interview: if the tooling were withdrawn tomorrow, what would break. That question surfaces dependencies respondents do not otherwise articulate, because it requires them to reason about the current arrangement rather than to describe it.

One question is deliberately absent. The protocol never asks whether artificial intelligence has reduced junior hiring, because the question supplies a socially available answer and respondents will reach for it, whether or not it describes their firm. The topic is reached through task-level description, and the inference is left to the respondent to draw or decline to draw.

3.4 Phase 3: critical incident technique

Flanagan's (1954) technique is underused in talent management research, and its value here is specific rather than general. Respondents asked about progression tend to narrate policy, because policy is what they have been trained to articulate and what they are accountable for. Asked instead to recount a particular occasion, they narrate events, and events resist tidying. Each respondent is asked for two incidents: one in which a junior engineer succeeded unexpectedly, and one in which the progression process failed. Incidents are recorded against a fixed frame covering antecedent conditions, actions taken, outcome, and the respondent's own attribution, with the attribution kept separate from the account so that the two can be analysed independently.

The target is 60 to 70 usable incidents across the interview set. Classification of incidents into categories is a codable step with reportable reliability, which gives the study a second point at which coding decisions are exposed to independent check. Incidents also anchor abstract claims to specific events, and a claim about progression that cannot be attached to any recounted incident is worth treating with suspicion in the analysis.

3.5 Coding, saturation, and reliability

Coding follows the Gioia methodology (Gioia, Corley, and Hamilton, 2013), which is appropriate where the aim is to build theory that retains informant terms while reaching theoretical abstraction. Its procedural lineage runs through grounded theory, and the coding practice here draws on the systematic variant of Strauss and Corbin (1998) rather than the constructivist one of Charmaz (2014), consistent with the codebook position taken below on reliability. First-order coding stays close to informant language and retains in-vivo codes, including the terms firms have coined for the

arrangements they are improvising. Second-order coding asks what theoretical concept a set of first-order concepts jointly indicates. Aggregate dimensions are the constructs that enter the model in Figure 1, which means the data structure and the propositional model must be read against each other rather than reported separately. Coding is conducted in NVivo or ATLAS.ti with the audit trail retained and exported as an appendix. Figure 3 shows the intended form of the data structure.




Figure 3. Data structure showing the progression from informant-centric first-order concepts through researcher-centric second-order themes to aggregate theoretical dimensions. First-order entries are illustrative of the form the data structure should take once fieldwork is complete; they are not findings.

Saturation is declared against a criterion set in advance, because retrospective assertions of saturation are unfalsifiable and reviewers have become appropriately impatient with them. The operationalisation adopted here follows Guest, Bunce and Johnson (2006): fewer than five percent new codes across three consecutive interviews. New codes per interview and the cumulative distinct code count are tracked throughout Phase 2 and reported in the form shown in Figure 4, with the reserve interviews released if the criterion is not met within the planned set.




Figure 4. Saturation monitoring. Bars show new codes generated per interview; the line shows cumulative distinct codes. The dashed marker indicates the point at which the pre-declared criterion is met. Values shown are illustrative of the reporting format, not data.

A second coder independently codes a random 25 percent subset. Cohen's kappa is reported for the categorical classification of critical incidents and Krippendorff's alpha for the thematic coding, the latter because alpha accommodates more than two coders and handles missing data (Krippendorff, 2018). The disagreement resolution procedure is reported alongside the coefficients rather than in place of them, since a coefficient without a resolution procedure says little about how the final coding was arrived at.

This choice carries a commitment that should be acknowledged rather than passed over. Reflexive thematic analysis in the Braun and Clarke (2006) tradition holds that intercoder reliability is conceptually inappropriate, on the grounds that coding is an interpretive act rather than the measurement of a fixed reality, and that agreement between coders demonstrates shared training rather than accuracy. The objection has force. This study reports reliability statistics and therefore positions itself as codebook thematic analysis rather than reflexive thematic analysis, accepting the associated constraint that the coding frame is treated as stable enough for two analysts to apply. Studies that report kappa while describing their method as reflexive are inconsistent on a point that reviewers in qualitative methods notice.

3.6 Phase 4: fuzzy-set qualitative comparative analysis

Phase 4 treats each firm as a configuration of conditions and asks which combinations are sufficient for capability depth. The technique suits RQ3 because the underlying argument is configurational: no single condition is expected to produce capability depth alone, and a firm with a wide entry portal but no structured review apprenticeship is expected to differ from one with both, in a way that an additive specification would not capture. Calibration follows the direct method with three anchors per condition, and truth-table analysis proceeds against a consistency threshold of 0.80. Figure 5 shows both steps in the form in which they will be reported.




Figure 5. (a) Direct calibration of a raw condition into fuzzy-set membership using three anchors. (b) Truth table format with filled and open circles denoting presence and absence of conditions, and consistency scores against a 0.80 threshold. Both panels illustrate the analytical format; the values are not results.

Five conditions and one outcome enter the analysis, with the operationalisation and anchors in Table 3. The anchors are justified theoretically rather than derived from the sample distribution, which is the step most frequently criticised in published fuzzy-set work and the one most often left implicit. The 20 percent full-in anchor for entry portal openness, for instance, reflects the graduate intake share historically required to sustain a pyramid structure in IT services, not the upper quartile of whatever distribution the sample happens to produce.

Table 3. Conditions, calibration anchors, and data source

Condition	Operationalisation	Calibration anchors	Source phase
Digital transformation intensity	Share of engineering workflow with tool-assisted execution, coded from interviews	Full in: pervasive; crossover: selective; full out: pilot only	Phase 2
Competence re-specification	Competency framework revised on judgement criteria	Full in: rewritten; crossover: amended; full out: unchanged	Phase 1
Entry portal openness	Graduate intake as share of total technical hiring	Full in: >20%; crossover: 10%; full out: <3%	Phases 1 and 2
Structured review apprenticeship	Formal assignment of AI-output review to juniors with mentorship	Full in: formalised; crossover: informal; full out: absent	Phases 2 and 3
Employment perimeter breadth	Share of technical workforce on direct employment	Full in: >85%; crossover: 60%; full out: <35%	Phase 1
Outcome: capability depth	Internal fill rate for senior technical roles	Full in: >70%; crossover: 45%; full out: <20%	Phase 1

Reporting follows three requirements. Necessity analysis is conducted before sufficiency and reported even where no condition proves necessary, since a null necessity result is informative about the configurational claim. Both consistency

and coverage are reported for every solution term, and the parsimonious and intermediate solutions are distinguished rather than merged. With 24 cases and five conditions, limited diversity is unavoidable, and the treatment of counterfactuals is stated explicitly rather than left to software defaults.

One consequence of the design arithmetic should be recorded here rather than discovered during analysis. Two of the five conditions, digital transformation intensity and structured review apprenticeship, are coded from Phase 2 and Phase 3 material, which means they can only be assigned for firms in the interview set. Reaching 24 firm-level cases therefore requires extending Phase 2 beyond the 30-interview core with shorter single-informant interviews in additional firms, and those cases will rest on thinner evidence than the fully stratified ones. Case-level evidence density should be reported alongside the truth table so that readers can see which configurations are carried by well-evidenced cases and which are not.

3.7 Trustworthiness

Trustworthiness procedures are mapped to the four criteria of Lincoln and Guba (1985) and reported as executed, with deviations noted, rather than listed as intentions. Credibility rests on triangulation across the three role strata and across document and interview evidence, together with member checking of the second-order themes with a subset of respondents. Transferability rests on thick description of firm type, size band, and transformation context, which is what allows a reader to judge applicability to another setting. Dependability rests on the retained coding audit trail and the pre-declared saturation criterion. Confirmability rests on a reflexive record of the researcher's prior assumptions, which matters in this study because the propositions predict a specific direction and confirmation is the easier finding to reach. Figure 6 maps procedures to criteria.




Figure 6. Trustworthiness procedures mapped to the four criteria of Lincoln and Guba (1985). Each procedure should be reported as executed, with deviations noted, rather than listed as intended.

4. Conclusion

The argument set out here can be wrong in several informative ways, which is the main reason for specifying it before fieldwork. If firms are found to have narrowed the entry portal while leaving progression policy unrevised, the mechanism operates roughly as proposed. If they are found to have revised progression policy and preserved the entry portal by assigning junior engineers to structured review of generated output, the mechanism is real but is being managed, and the more useful question becomes what distinguishes the firms managing it, which is precisely the question Phase 4 is built to answer. If entry portals turn out to have narrowed uniformly regardless of transformation intensity, the financing explanation is stronger than the account offered here, and the paper should report that.

What the design commits to is making all three outcomes reportable. That commitment is what the disconfirming-evidence column in Table 4, the within-firm and across-stratum comparison, and the pre-declared saturation criterion are for. It is also why the propositions are stated as directional claims about observable artefacts, since a proposition that cannot fail against a competency framework or a graduate intake share is not doing analytical work. The entry-portal

problem, if it exists in the form described, has a structure that no individual firm can solve on its own terms, and identifying that structure is a more useful contribution than a list of practices firms should adopt.

References

- Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120.
- Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis, with special reference to education*. University of Chicago Press.
- Boudreau, J. W., & Ramstad, P. M. (2005). Talentship, talent segmentation, and sustainability: A new HR decision science paradigm for a new strategy definition. *Human Resource Management*, 44(2), 129–136.
- Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27–40.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.
- Cappelli, P. (2008). Talent management for the twenty-first century. *Harvard Business Review*, 86(3), 74–81.
- Charmaz, K. (2014). *Constructing grounded theory (2nd ed.)*. Sage.
- Collings, D. G., & Mellahi, K. (2009). Strategic talent management: A review and research agenda. *Human Resource Management Review*, 19(4), 304–313.
- Flanagan, J. C. (1954). The critical incident technique. *Psychological Bulletin*, 51(4), 327–358.
- Gallardo-Gallardo, E., Dries, N., & Gonzalez-Cruz, T. F. (2013). What is the meaning of talent in the world of work? *Human Resource Management Review*, 23(4), 290–300.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31.
- Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59–82.
- Krippendorff, K. (2018). *Content analysis: An introduction to its methodology (4th ed.)*. Sage.
- Lewis, R. E., & Heckman, R. J. (2006). Talent management: A critical review. *Human Resource Management Review*, 16(2), 139–154.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage.
- Mitchell, T. R., Holtom, B. C., Lee, T. W., Sablinski, C. J., & Erez, M. (2001). Why people stay: Using job embeddedness to predict voluntary turnover. *Academy of Management Journal*, 44(6), 1102–1121.
- Ragin, C. C. (2008). *Redesigning social inquiry: Fuzzy sets and beyond*. University of Chicago Press.
- Rousseau, D. M. (1995). *Psychological contracts in organizations: Understanding written and unwritten agreements*. Sage.
- Sahoo, D. R. (2017a). Creative Accounting: A Critical Study. *International Journal of Multifaceted and Multilingual Studies*, 4(IX).
- Sahoo, D. R. (2017b). Emotional Intelligence: Managing Stress and Anxiety at work place. *International Journal of Multifaceted and Multilingual Studies*, 4(IX).
- Sahoo, D. R. (2018). Corporate Governance: Need of the Hour (Companies Act 2013). *International Journal of Management Studies (IJMS)*, 4(ii).
- Sahoo, D. R., & Deepak, S. V. (2024a). Exploring the Ethical Perspectives of Sustainable Finance: A Research Study. In *Advances in Business Strategy and Competitive Advantage* (pp. 119–142). IGI Global.
- Sahoo, D. R., & Deepak, S. V. (2024b). Exploring the Nexus of Sustainable Finance: ESG Ethical Framework Practices. In *Ethical Quandaries in Business Practices: Exploring Morality and Social Responsibility* (pp. 253–280). IGI Global.
- Sahoo, D. R., & Jagtap, V. K. (2017). IND AS 109: Its Impact on Financial Statements. *International Research Journal of Multidisciplinary Studies*, 3(12), 1–5.
- Sahoo, V. D., & Sahoo, D. R. (2025a). Promoting Inclusivity: Corporate Social Responsibility in Game-Based Learning and Professional Employment. In *Game-Based Education Approaches to Inclusive Business Management* (pp. 101–126). IGI Global Scientific Publishing.
- Sahoo, V. D., & Sahoo, D. R. (2025b). *Prospective of Blockchain in Derivative Markets: An Empirical Review*. *Innovations in Blockchain-Powered Intelligence and Cognitive Internet of Things (CIoT)*, 229–252.

- Shende, M. A., & Sahoo, D. R. (2021). Talent Management Critical Implications and Strategic Approach Reference to Employer and Employees Prospective. *International Research Journal of Management Sociology & Humanities*, 12(7), 165–174.
- Sparrow, P. R., & Makram, H. (2015). What is the value of talent management? Building value-driven processes within a talent management architecture. *Human Resource Management Review*, 25(3), 249–263.
- Strauss, A., & Corbin, J. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory (2nd ed.)*. Sage.
- Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350.
- Vaiman, V., Scullion, H., & Collings, D. (2012). Talent management decision making. *Management Decision*, 50(5), 925–941.
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *Journal of Strategic Information Systems*, 28(2), 118–144.
- Warner, K. S. R., & Wäger, M. (2019). Building dynamic capabilities for digital transformation: An ongoing process of strategic renewal. *Long Range Planning*, 52(3), 326–349.
- Weil, D. (2014). *The fissured workplace: Why work became so bad for so many and what can be done to improve it*. Harvard University Press.